

Predicting the Occurrence of Broken Rails: A Quantitative Approach

C. Tyler Dick
Graduate Research Assistant
Railroad Engineering
University of Illinois at Urbana-Champaign

Current address:
HDR Engineering, Inc.
512 Main Street, Suite 600
Fort Worth, TX 76102
Tel: (817) 810-9132 Fax: (817) 810-0772
cdick@hdrinc.com

Christopher P.L. Barkan
Associate Professor
Director - Railroad Engineering Program
University of Illinois at Urbana-Champaign
1201 Newmark Civil Engineering Laboratory
205 N. Mathews Ave., Urbana, IL, USA 61801
Tel: (217) 244-6338 Fax: (217) 333-1924
cbarkan@uiuc.edu

Edward Chapman
Director - Hazardous Materials
Burlington Northern & Santa Fe Railway
2600 Lou Menk Drive
Fort Worth, TX, USA 76131
Tel: (817) 352-1954
ed.chapman@bnsf.com

Mark P. Stehly
Assistant Vice President - Environment and Research & Development
Burlington Northern & Santa Fe Railway
2600 Lou Menk Drive
Fort Worth, TX, USA 76131
Tel: (817) 352-1907 Fax: (817) 352-7225
mark.stehly@bnsf.com

ABSTRACT

A multivariate statistical model was developed to improve the ability to predict the occurrence of broken rails, (aka "service failures"). Broken rails are the leading cause of major accidents on US railroads and service failures are a frequent cause of delays. The model uses a combination of engineering and traffic data commonly recorded by major railroads. The Service Failure Prediction Model (SFPM) enables railroads to identify the conditions that are predictive of the occurrence of broken rails. This can help them allocate inspection, detection and preventive resources more efficiently, thereby enhancing safety, reducing risk and service interruptions, and potentially extracting more use from their rail assets.

INTRODUCTION

Derailments caused by broken rails have been a safety concern of the railway industry for over a century (1,2). Improvements in rail manufacturing, inspection and rail defect detection have greatly reduced the incidence of broken rails; however, they remain a frequent cause of service interruptions and one of the leading causes of derailments. Improving the ability to predict where broken rails are likely to occur has both economic and safety benefits because it would enable more effective allocation of resources to detect and prevent broken rails (3,4). Previous models have not had the benefit of the availability of large, readily accessible, databases containing extensive detail on many key parameters likely to effect the occurrence of broken rails (5,6,7).

Definition of Severe Derailments

We were interested in identifying the derailment causes most likely to lead to a severe accident. Severe derailments were defined as those in which a large number of cars derailed at speed. Such accidents will, in general, have the greatest potential for harm to persons, property, equipment and track. Furthermore, analyses of Federal Railroad Administration (FRA) accident data showed that accidents with these characteristics were strongly correlated with release of hazardous materials if they were present in the derailed portion of the train (8). Consequently, for both safety and economic reasons, prioritizing attention on these types of derailments was of interest.

Derailment Severity-Frequency Analysis

To determine the causes of accidents most likely to lead to severe derailments, we conducted a simple risk analysis using FRA data (9) for 3,504 mainline derailments that occurred during the five-year interval 1994 - 1998. The FRA reporting system requires identification of a primary cause (and other contributing causes if applicable). Data on accident causes are grouped hierarchically by FRA and we used data at the FRA subcause level (the second-highest level of aggregation). We calculated the average number of cars derailed in accidents attributed to each subcause and plotted these against the frequency of derailments caused by the same subcause (Figure 1).

Figure 1 is divided into four quadrants by vertical and horizontal lines that represent the average value of the two variables with respect to the X and Y axes, respectively. The vertical line represents the average frequency of accidents for all recorded causes combined, and the horizontal line is the average number of cars derailed due to each cause. Causes above or below these lines are, by definition, above or below average for the respective axis.

The causes in the upper right quadrant are most interesting and pose the greatest risk because they are both more frequent and more severe than average. The FRA cause code "Rail and joint bars" is clearly the highest consequence, high frequency cause of accidents. Further analysis of the FRA accidents attributed to this cause reveals that most were due to broken rails. Based on these results we undertook a more detailed analysis of the factors contributing to the occurrence of broken rails (8). Several recent hazardous materials accidents have underscored the importance of this particular aspect of the risk from broken rails.

Broken Rails

Most broken rails do not result in derailments. Instead the break is detected, usually by the track circuit system or by track inspectors, and repaired (on several North American railroads these detected broken rails are referred to as "service failures"). Broken rail derailments appear to be correlated with the occurrence of service failures (8). Therefore predicting the occurrence of service failures has a potential safety benefit because it could enable railroads to allocate broken rail prevention measures, detection technology or inspection efforts, more effectively (3,4,6). Furthermore, understanding the factors correlated with service failure occurrence could help identify contributing causal factors, thereby enabling better preventive measures. The objective of this research was to develop a probabilistic model to predict the circumstances most likely to lead to the occurrence of a service failure.

Model Form and Data Set

Ideally, the model we developed would enable the user to input values for the relevant parameters at a specific location on the railroad and determine a measure of the probability of a service failure there. The output of the model is an index value between 0 and 1, with 0 indicating the lowest probability of service failure and 1 representing the highest. Since a probability is the desired output and there are only two possible outcomes, service failure or no service failure, at each location, the model can be constructed as a discrete choice model.

A discrete choice model, such as the logit model, fits an appropriate equation to the data and uses this equation to score each location relative to a threshold value, above which failure is predicted to occur. The logit model then uses a logistic distribution to consider the uncertainty and error in the estimated score and threshold value, and determine the probability that the score is above the threshold value. The calculated probability is then used as an estimate of the service failure probability at that particular location.

In order to fit a discrete choice logit model, two sets of data were required; one to characterize locations where service failures occurred and a second set of data to characterize locations where service failures had not occurred. Development of these data began with information the Burlington Northern and Santa Fe Railway had developed containing detailed information on the date, location and type of 1,903 service failures that occurred over a two-year interval. These data were supplemented with engineering and operational data pertaining to each service failure location. A new dependent variable was created and assigned the value "1" for each of these records signifying that a service failure had occurred there.

The second set of data was created with records for locations where no service failure occurred during the same interval. An approximately equal-sized set of data was developed by selecting a random sample of locations from the railroad and assembling identical information as had been developed for the service failure locations. The dependent variable for these records was assigned a value of "0".

Ultimately, we developed a test database comprising 3,676 records with complete service failure and descriptive parameter information. Based on a univariate analysis of the service failure data and review of literature on the circumstances of rail defect growth and broken rail occurrence (7, 10, 11), track structure and dynamic effects (12, 13, 14), and the fracture mechanics of rail (5, 15) the following parameters were selected to be considered in the multivariate service failure model:

- Rail Age
- Rail Weight
- Degree of curve
- Speed
- Average Tons Per Car
- Average Dynamic Tons Per Car
- Percent Grade
- Annual Gross Tonnage
- Annual Wheel Passes
- Insulated Joints
- Mainline Turnouts

All of the parameters are continuous variables except the last two, insulated joints and mainline turnouts, which are both discrete. These were assigned a value of "1" if present at a location, and "0" if not.

Model Development

The service failure probability model was developed using Statistical Analysis Software (SAS) and the LOGISTIC procedure. The LOGISTIC procedure fit a discrete choice logit model to the test database. Stepwise regression was used to determine the most relevant parameters and combinations of parameters (two-factor interaction terms) for inclusion in the model. The stepwise regression procedure uses an iterative process to select variables on the basis of their ability to explain the variance in the input data. The model conducts a "goodness-of-fit" test for each step and adds or subtracts variables or combinations of variables, until the addition of another parameter does not significantly improve the fit. At this point the last version of the model is considered the "best" and the resultant parameters, coefficients, and functional relationships comprise the final model.

Retrospective and Prospective Models

Development of the service failure model was a two step process. First the model was fit to the test database described above. Recall that this database comprised 3,676 locations, approximately half of which experienced a service failure during the two-year period encompassed, and the other half were a random sample of locations that did not. Because the model is making predictions about the past, we termed it the "retrospective model". This version of the model is used primarily to assess the accuracy of the model's predictions with respect to the test database.

The second step of the process is development of a "prospective model". It is modified from the retrospective model by adjusting a constant term to reflect the actual

average service failure probability over whatever portion of a railroad system is of interest. Once this adjustment is made, the prospective model can be used to calculate the annual probability of a service failure at particular locations, or along any portion of track that is of interest.

Retrospective Service Failure Model. The retrospective service failure probability model was developed using the LOGISTIC procedure:

$$p_{SF2} = e^U / (1+e^U)$$

Where:

p_{SF2} = probability that a service failure occurred at a particular point during the study interval

$$U = Z + Y$$

Z = -4.569, model specific constant, (discussed below)

$$Y = 0.059A + 0.025AC - 0.00008A^2C^2 + 5.101T/S + 217.9W/S - 3861.6W^2/S^2 + 0.897(2N-1) - 1.108P/S$$

A = rail age in years

C = degree of curvature (= 0 for tangent)

T = annual traffic in million gross tons (MGT)

S = rail weight in pounds

W = 4T/L = annual number of wheel passes (millions)

P = L(1 + V/100) = dynamic wheel load

N = 1 if at turnout, 0 if not at turnout

L = tons per car

V = track speed

The fitted model includes a model-specific constant or intercept term, Z, that is related to the average service failure probability. Recall that the retrospective model is fit to a data set in which approximately half of the records were for locations that experienced service failures. The average service failure probability on an actual system would be far lower, so this term would be adjusted to reflect this (see discussion of prospective model below).

Interpretation of Model Terms. The service failure probability model contains terms that describe different effects and relationships between service failure probability, infrastructure characteristics and traffic characteristics.

The first term in the model, $0.059A$, reflects the effect of rail age. As rail age increases, service failure probability increases. This result is consistent with extensive industry experience. Older rail is likely to have carried more tonnage, experienced more thermal stress cycles and may have been manufactured using processes that allowed more flaws in the rail. A recent study of rail failures on Railtrack in Great Britain (16) supports the importance of this parameter.

The second and third terms in the model, $0.025AC - 0.00008A^2C^2$, reflect the interaction between rail age and degree of curve. As either rail age or degree of curve increases, service failure probability is predicted to increase. Since the interaction between rail age and curvature is multiplicative, the model indicates that in terms of

service failure probability, higher degree (sharper) curves are more sensitive to the effects of rail age, and vice versa.

The fourth term in the model, $5.101T/S$, reflects the effect of annual traffic (MGT) normalized by rail weight. As annual gross tonnage increases, service failure probability increases. However, the form of the interaction with rail weight indicates that the increase in service failure probability associated with a unit increase in annual traffic is greater on segments of track with relatively light rail.

The fifth and sixth terms in the model, $217.9W/S$ and $3861.6 W^2/S^2$, describe the effect of annual wheel passes or load cycles normalized by rail weight. Service failure probability increases as the number of wheel passes or load cycles increases. However just as with gross tonnage, the increase in service failure probability associated with a unit increase in the annual number of wheel passes is greater on segments of track with relatively light rail. This is probably due to the fact that lighter rail experiences more stress under a given load than heavier rail. Thus, the amount of crack growth per fatigue cycle is greater in lighter rail than heavier rail.

It is interesting that the model includes terms that describe annual traffic in terms of gross tonnage and the number of wheel passes. The relationship between annual traffic and service failure probability is a function of both the total amount of load applied to a section of rail and the number of times the load is applied. This relationship is consistent with fracture mechanics models of fatigue crack growth in rails that depend on both the applied stress and the number of load cycles (15).

The seventh term in the model, $0.897(2N-1)$, describes the effects of mainline turnouts. Since $N = 1$ near a turnout, the presence of a turnout increases the probability of a service failure. There are several possible explanations related to inferences about rail stress. Turnouts may tend to anchor the track structure thereby causing greater thermal stress cycling as the nearby rail expands and contracts. Also, to the extent that turnouts tend to be associated with locations where trains slow down, stop or start, rails in these locations may tend to experience more traction-induced stresses.

The final term in the model, $-1.108P/S$, describes the effect of dynamic load on service failure probability. The term is negative indicating that as dynamic load increases, service failure probability decreases. This is an unexpected result and is the opposite of what was suggested by a single variable analysis conducted prior to development of the multi-variate model. However, the relative effect of this term is weak. For example at an annual tonnage level of 50 MGT, on 136 lb. rail, in tangent track, varying the annual wheel passes between the highest and lowest possible values changes p_{SF2} by approximately 0.17. Under the same conditions, varying the dynamic load term between its extreme values only changes p_{SF2} by 0.03. In the stepwise regression this term was the final term added to the model (Table 1) and has the least predictive ability of the included terms (as indicated by the low chi-squared value). Thus we do not think that this term represents an actual physical relationship. The regression model development procedure may have included this term in the model to capture additional, unexplained variance resulting from various effects, and possibly to balance over-predictions of service failure probability caused by the linear nature of other effects in the model.

Table 1 also indicates that during the stepwise regression process, an interaction term between rail age and annual gross tonnage was initially included in the model. By multiplying rail age by annual gross tonnage, the term estimated the effect of cumulative tonnage. Although the cumulative tonnage effect was initially significant, as more detailed terms describing the effects of rail age, turnouts and curvature were added to the model, the cumulative tonnage effect became less significant and was finally removed from the model. Thus, the variance in service failure probability that was initially explained by cumulative tonnage in a model with two terms could be better explained by a model with more terms and a combination of effects involving other variables. We would like to have included actual accumulated tonnage, but this variable was not consistently available on a system-wide basis.

It is also interesting which parameters did not appear in the final model. The effects of grade, speed, average wheel load and insulated joints were not found to significantly improve the predictive ability of the model and were not included. Some variables that we would have liked to consider were, rail steel type, rail surface roughness, neutral temperature, and actual temperature at the time of the break, but these data were not available.

Retrospective Service Failure Model Performance. We used two methods to evaluate the ability of the retrospective model to predict locations where service failures occurred. The first calculates a goodness of fit statistic for the model based on the service failure probability (p_{SF2}) computed for each of the records in the input data. If the model completely accounted for all of the sources of variance, one would expect $p_{SF2} = 1$ at all of the service failure locations and $p_{SF2} = 0$ at all of the locations where service failures did not occur. In this case the summation of p_{SF2} over all service failure locations should equal the total number of service failures and the summation of $1 - p_{SF2}$ over all locations where service failures did not occur should equal the total number of locations where service failures did not occur. It is highly unlikely that all sources of variance will have been accounted for by any statistical model. Therefore, when the summations are computed for actual values of p_{SF2} , they will correctly account for only a percentage of the total. This percentage reflects the "goodness of fit" or the amount of variance explained by the retrospective model (17). Using this approach, the goodness-of-fit statistic is calculated using the expression below, where n_{sf} is the actual number of locations where service failures occurred, and n_{nosf} is the number of locations where they did not.

$$\begin{aligned} \text{Goodness of fit} &= \left(\sum_{n_{sf}} p_{SF2} + \sum_{n_{nosf}} 1 - p_{SF2} \right) / (n_{sf} + n_{nosf}) \\ &= (1,507 + 1,462) / (1,861 + 1,815) \\ &= 0.808 \end{aligned}$$

Based on this analysis, the retrospective model accounted for 80.8 percent of the variance in the service failure data.

The second method we used to evaluate the performance of the model is to compare the value of p_{SF2} to the event that actually occurred at that location. The decision criteria, or threshold value, for service failure prediction was a p_{SF2} value of 0.5. If p_{SF2} was less than 0.5, it was classified as predicting no failure and if it was greater than 0.5 it was classified as predicting a service failure. 87.4 percent of these predictions were correct (Table 2). Of the incorrect predictions, there were twice as many false positives than missed service failures. This indicates that the model is somewhat conservative as it is more likely to provide a false positive than miss a service failure. The decision criteria of 0.5 could be adjusted by users of the model to make results more or less conservative (δ).

These two evaluations indicate that the model had a reasonably high level of accuracy in predicting the occurrence of service failures in the database from which it was developed. The next steps in assessing the model's accuracy would be to apply it to data for a subsequent time period on the same railroad, and to apply it to data from a different railroad.

Prospective Service Failure Model. As explained above, in order to use the model to predict the annual probability of a service failure at a particular location, the retrospective model must be transformed into a prospective model. This transformation is accomplished by adjusting the value of the model specific constant, Z , to reflect the average service failure probability across the entire system of interest. There were 1,861 service failures in the test database over a two-year period for which complete records were available. The probability that one of these service failures falls into any given segment of track is a function of the length of the segment. To capture as much detail as possible, and to avoid the use of average values over a segment that may introduce additional variance, the segments should be kept relatively short. The maximum resolution in the data available for most of the parameters of interest was 0.01 miles (16.09 m). The total system represented by the database was approximately 23,750 miles of mainline. Thus, there were 2,375,000 segments 0.01 miles in length. Given this value, the average probability that a service failure is found in any one of those segments over a two-year period is approximately 0.00078. This probability can be converted into a new model-specific constant, Z^* , through the use of the log-odds operator (1δ):

$$\begin{aligned} Z^* &= Z + \ln [p_{SFavg} / (1 - p_{SFavg})] \\ &= -4.569 + \ln [0.00078 / (1 - 0.00078)] \\ &= -11.763 \end{aligned}$$

This new model specific constant, Z^* , adjusts the scale of the probability calculated by the prospective service failure model so that the model predicts service failures at a rate that is comparable to the actual observed rate.

The retrospective model described above calculated the probability of a service failure over a two-year period. This can be converted to an annual probability simply by dividing by two when transforming the U-score into a probability. Once these two adjustments are made, the annual service failure probability for any 0.01 mile segment

on the system can be calculated with the prospective service failure model. The prospective service failure probability model has the following form:

$$p_{SF} = e^U / [2(1+e^U)]$$

Where:

p_{SF} = annual probability of a service failure in the 0.01-mile segment of interest

$$U = Z^* + Y$$

$Z^* = -11.763$, prospective model specific constant, described above
(all other terms and variables are the same as defined previously)

Service Failure Probability and Expected Service Failures per Mile

Cursory review of the annual service failure probabilities calculated by the prospective model might suggest they are too low. However, the probability is based on a segment of track that is only 0.01 miles in length. The calculated probability is approximately equal to the expected number of service failures per year in that 0.01 mile segment. Annual service failures per mile is a metric more typically used by North American railroads, so it is useful to calculate a per-mile rate by multiplying p_{SF} by 100.

$$SF/MI/YR = 100e^U / [2(1+e^U)]$$

Where:

SF/MI/YR = expected service failure rate on segment of interest (service failures per mile per year)

This rate can be applied to a segment of track of any length as long as the values of the parameters in the service failure model remain constant along the section of the track. A service failure rate of 2 SF/MI/YR indicates that for every mile of track for which the rate applies, two service failures are expected to occur. If the track section to which this rate applied is 0.5 miles in length, then one service failure is expected along this length and if the section is two miles in length, four service failures are expected along the two-mile length. Note that in all three cases the service failure rate, 2 SF/MI/YR, is the same but the number of service failures expected in a section of track is a linear function of the length of the section. The number of service failures expected in a given section of track where the service failure rate is constant can be calculated by multiplying the length of the section by the service failure rate.

Example of Service Failure Model Application

The following example illustrates how the prospective service failure prediction model (SFPM) can be used to obtain a measure of service failure probability and rate. A hypothetical 1.5-mile, single track portion of a railroad mainline is illustrated in Figure 2 and the relevant parameters are presented in Table 3. The segment has been broken into several sub-segments over which the input parameters are constant.

Some of the rail in the segment of interest is 47 years old and weighs 132 pounds per yard. The remaining rail is five years old and weighs 136 pounds per yard. Mainline turnouts are located at mile zero and also at mile 0.7 where another mainline connects to the line under study. A one-degree curve is located between mile 0.25 and mile 0.45. Track speed on the segment is 50 miles per hour. The annual traffic is 80 million gross tons between mile 0.0 and mile 0.7. At mile 0.7, 40 million gross tons is routed on the connecting mainline with the remaining 40 million gross tons being routed on the segment under consideration between mile 0.7 and 1.5. The average gross rail load is 100 tons in the eastbound direction and 80 tons in the westbound direction for a maximum of 100 tons. The dynamic load computes to 150 tons per car, and the annual traffic of 80 MGT and 100-ton average per car results in an estimated 3.2 million wheel passes.

The U-score was calculated for each portion of the segment of interest and then transformed into an estimate of service failure rate. The estimated service failure rate (service failures per mile per year) for each sub-segment is summarized in Table 4 and presented graphically in Figure 3. Multiplying the service failure rate on each sub-segment by the actual length of each sub-segment provides an estimate of the expected number of service failures on an annual basis per mile in that sub-segment. Summing all of the individual sub-segment values provides an estimate of the expected number of service failures per year on all 1.5 miles of the segment of interest. In this case, the expected number of service failures for the segment is 0.316.

The service failure profile in Figure 3 highlights how interactions between the various parameters affects service failure rate. Between mile 0 and 0.1, the rail is relatively old and a turnout is present. The combination of these two factors results in a relatively high predicted service failure rate. At mile 0.1, the service failure rate drops as the rail is no longer close enough to the turnout to be subject to its effects. Between mile 0.1 and 0.25, the track is tangent but the old rail produces a higher service failure rate than on the segment between mile 0.45 and 0.6 where the track is tangent but the rail is relatively new. This difference in service failure rate illustrates the importance of rail age. Under the traffic conditions in this example, the age difference of 42 years results in a service failure rate that is 16 times higher on the older section of rail. At mile 0.25, the track transitions from tangent to a one-degree curve and the service failure rate increases approximately three times. When compared to mile 0.45, where the new rail transitions from curve to tangent and the service failure rate only increases by a factor of 1.5, the increase in service failure rate at mile 0.25 is large. This is due to the interaction of rail age and curvature that makes the old rail on this sub-section of track sensitive to curvature. At mile 0.3, the rail on the one degree curve changes from rail that is 47 years old to rail that is 5 years old. The model suggests that newer rail is less sensitive to curvature, so the service failure rate drops from 0.86 to 0.03 service failures per mile per year. Since there is one half the traffic between mile 0.7 and 1.5 than there is between mile 0.0 and 0.7, the service failure rate is also correspondingly lower.

CONCLUSIONS

A simple risk analysis showed that broken rails are the leading cause of severe accidents as measured by number of cars derailed. Improved detection and prevention

of broken rails has important potential safety and economic benefits. Furthermore, there are service quality and reliability benefits if the incidence of broken rails can be reduced. Improving the ability to predict the conditions that can lead to broken rails can help railroads allocate inspection, detection and preventive resources more efficiently, thereby enhancing safety and reducing service interruptions due to broken rails.

We developed a statistical model that provides probabilistic estimates of the likelihood of service failure occurrence based on engineering and operational input parameters. Although further validation needs to be conducted, the service failure prediction model described here shows promise of being able to provide improved ability to predict the occurrence of broken rails. If the requisite data for a railway system can be systematically developed in a consistent, easily accessed, electronic format, the model described in this paper can be applied to any portion of the system to generate probabilistic estimates of service failure probability. If the data include appropriate geographical information, then the service failure model presented here could be incorporated into a geographic information system that would generate service failure and broken rail derailment profiles automatically from railway databases.

Previous models have been based on a combination of fundamental principles and a limited number of variables available on rail in service. The information technology and computer revolution has led to the availability of much larger, comprehensive databases and made feasible the use of more powerful statistical tools. The research described here would not have been feasible 10 years ago. The results of these analyses, coupled with sophisticated graphical output, can improve managers' access to information and enhance the quality and pace of decision-making. The potential benefit of the approach is greater precision in predicting the occurrence of broken rails, along with wider availability and enhanced interpretation of the results. This is important as railroads strive to improve safety, while at the same time make more efficient use of their resources and extract more value from assets such as rail.

ACKNOWLEDGMENTS

Douglas Simpson and Todd Treichel provided helpful assistance and review of the statistical methods used. Thanks also to Hank Lees, Tom Wright and Scott Staples who assisted us in obtaining the data needed for the analysis. Frederick Lawrence patiently shared his insight regarding the fracture mechanics of rail, and Kevin Sawley and David Davis also provided helpful discussion. The first two authors would like to express their gratitude to the Burlington Northern and Santa Fe Railway for its support of this research.

REFERENCES

1. Thompson, A.W. Service Performance of Railroad Rails, In: A.W. Thompson (*ed.*) *Symposium on Railroad History, Volume 2*, A.C. Kalmbach Library, Chattanooga, Tennessee, 1992.
2. Aldrich, M. The Peril of the Broken Rail: The Carriers, the Steel Companies, and Rail Technology, 1900-1945, *Technology and Culture*, Vol. 40, 1999, pp. 263-291.

3. Palese, J.W. and T.W. Wright. Risk Based Ultrasonic Rail Test Scheduling on Burlington Northern Santa Fe, In: *AREMA: Proceedings of the 2000 Annual Conference CD-ROM*, American Railway Engineering and Maintenance-of-Way Association, Landover, MD, 2000.
4. Palese, J.W. and A.M. Zarembski. BNSF Tests Risk-based Ultrasonic Detection, *Railway Track and Structures*, Vol. 97, No. 2, 2001, pp. 17-21.
5. Orringer, O. and M.W. Bush. Applying Modern Fracture Mechanics to Improve the Control of Rail Fatigue Defects in Track, *American Railway Engineering Association*, Volume 83, Bulletin 689, 1982, pp. 19-53.
6. Orringer, O., Y.H. Tang, D.Y. Jeong, and A.B. Perlman. *Risk/Benefit Assessment of Delayed Action Concept for Rail Inspection*, Office of Research and Development, Federal Railroad Administration, Washington, DC, 1999.
7. Lawrence, F.V., C.P.L. Barkan, and H.M. Reis. *Summary of a Workshop on New Technologies for the Detection of Broken Rail*, University of Illinois at Urbana-Champaign Railroad Engineering Program, Urbana, Illinois, 2000.
8. Dick, C.T. *Factors Affecting the Frequency and Location of Broken Railroad Rails and Broken Rail Derailments*. M.S. Thesis. University of Illinois at Urbana-Champaign, Urbana, IL, 2001.
9. Federal Railroad Administration, *Accident Incident Database 1994-2000, electronic database*, <http://safetydata.fra.dot.gov/officeofsafety/>, accessed May 2000.
10. Reiff, R.P. *Proceedings of Rail Defect and Broken Rail Defects Expanded Workshop July 1997*, Transportation Technology Center, Pueblo, CO, 1997.
11. Clark, R. Defective and Broken Rails - How They Occur, In: *F.V. Lawrence, C.P.L. Barkan, and H.M. Reis (eds.), Summary of a Workshop on New Technologies for the Detection of Broken Rail*, University of Illinois at Urbana-Champaign Railroad Engineering Program, Urbana, IL, 2000.
12. Hay, W.W. *Railroad Engineering*, Second Edition, Wiley & Sons, New York, NY, 1982.
13. Selig, E.T. and J.M. Waters. *Track Geotechnology and Substructure Management*, Thomas Telford Publications, London, 1994.
14. American Railway Engineering and Maintenance-of-Way Association. *Manual for Railway Engineering, CD-ROM*, 2000, pp. 16-10-9.
15. Lawrence, F.V. An Analytical Model for Broken Rail, In: *F.V. Lawrence, C.P.L. Barkan, and H.M. Reis (eds.), Summary of a Workshop on New Technologies for the Detection of Broken Rail*, University of Illinois at Urbana-Champaign Railroad Engineering Program, Urbana, Illinois, 2000.

16. Sawley, K. and R. Reiff. *Rail Failure Assessment for the Office of the Rail Regulator*, Transportation Technology Center, Association of American Railroads, Pueblo, CO, 2000.
17. Ben-Akiva, M. and S.R. Lerman. *Discrete Choice Analysis*, MIT Press, Cambridge, MA, 1985.
18. McCullagh, P. and J.A. Nelder. *Generalized Linear Models*, Second Edition, Chapman and Hall, New York, NY, 1989.

FIGURE 1. Frequency/severity graph of mainline derailments 1994-1998




FIGURE 2: Schematic of hypothetical section of mainline track

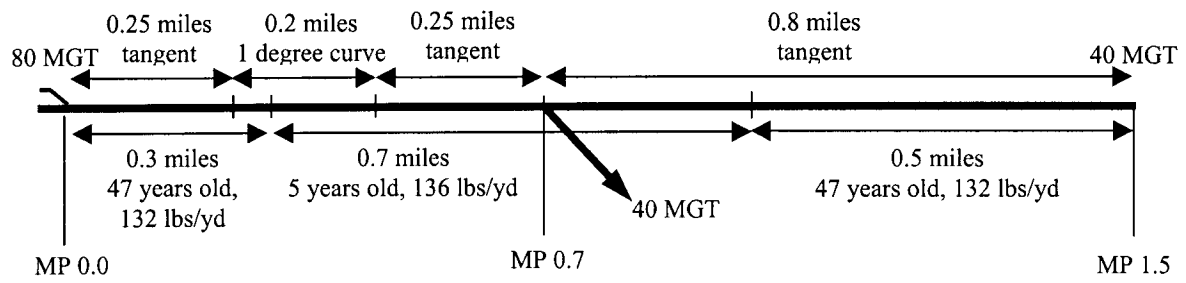


FIGURE 3. Graphical representation of service failure probability along a hypothetical track segment




TABLE 1: Model term selection order

Step	Term Added	Term Removed	Chi-Squared
1	Wheel Passes / Rail Weight	--	155
2	Annual Gross Tonnage x Rail Age	--	--
3	(Wheel Passes / Rail Weight) ²	--	202
4	Annual Gross Tonnage / Rail Weight	--	63
5	Rail Age	--	204
6	Turnout	--	41
7	Degree of Curve x Rail Age	--	47
8	(Degree of Curve x Rail Age) ²	--	33
9	Dynamic Load / Rail Weight	--	8
10	--	Annual Gross Tonnage x Rail Age	--

TABLE 2. Results of goodness-of-fit test for a threshold value of $p_{SF2}=0.5$

Model Prediction	Actual Event	Events	Percent of Total	Outcome
Service Failure ($p_{SF2} > 0.5$)	Service Failure	1,700	87.4	Correct
No Failure ($p_{SF2} < 0.5$)	No Failure	1,513		Prediction
Service Failure ($p_{SF2} > 0.5$)	No Failure	302	8.2	False Positive
No Failure ($p_{SF2} < 0.5$)	Service Failure	161	4.4	Missed Failure

TABLE 3: Input parameters for hypothetical section of mainline track

Start MP	End MP	Z	A (age)	C (degree)	T (MGT)	S (pounds)	W (million)	P (tons)	N (turnout)	U
0.00	0.10	-11.763	47	0	80	132	3.2	150	1	-3.25
0.10	0.25	-11.763	47	0	80	132	3.2	150	0	-5.04
0.25	0.30	-11.763	47	1	80	132	3.2	150	0	-4.04
0.30	0.45	-11.763	5	1	80	136	3.2	150	0	-7.47
0.45	0.60	-11.763	5	0	80	136	3.2	150	0	-7.60
0.60	0.70	-11.763	5	0	80	136	3.2	150	1	-5.80
0.70	0.80	-11.763	5	0	40	136	1.6	150	1	-8.26
0.80	1.00	-11.763	5	0	40	136	1.6	150	0	-10.05
1.00	1.50	-11.763	47	0	40	132	1.6	150	0	-7.53

TABLE 4. Service failure probabilities along hypothetical track section

Start MP	End MP	Length	U	SF/MI/YR	Expected SF
0.00	0.10	0.10	-3.25	1.866	0.187
0.10	0.25	0.15	-5.04	0.322	0.048
0.25	0.30	0.05	-4.04	0.865	0.043
0.30	0.45	0.15	-7.47	0.028	0.004
0.45	0.60	0.15	-7.60	0.025	0.004
0.60	0.70	0.10	-5.80	0.151	0.015
0.70	0.80	0.10	-8.26	0.013	0.001
0.80	1.00	0.20	-10.05	0.002	0.000
1.00	1.50	0.50	-7.53	0.027	0.014
Total (0.0 to 1.5)				0.211	0.316